

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant(s) : Edward K.Y. Jung et. al.
Application No. : 10/816,358
Filed : March 31, 2004
TITLE : AGGREGATING MOTE-ASSOCIATED DATA
Conf. No. : 1749
Examiner : Oleg Survillo
Art Unit : 2442
Docket No. : SE1-0013-US
Customer No. : 80118

Mail Stop Amendment
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

DECLARATION PURSUANT TO 37 CFR §1.132

1. My name is Richard T. Lord.
2. I have a degree in Electrical Engineering from Ohio University.
3. I have 45 years of technical experience in addition to my formal education.
4. I have held the following technical positions:

Instructor, Electrical Engineering, Ohio University

Application Engineer, General Electric Co.

Manager, Product Planning, General Electric Co.

Manager, Communication and Control Devices, General Electric Co.

Manager, Corporate Consulting Services, General Electric Co.

Manager, Product and Program Planning, Control Data Corp.

Director, Computer Peripherals, Centronics Data Computer Corp.

Manager, Telemetry and Control Systems, Siemens Energy & Automation, Inc.

Vice President, Energy Innovations, Dayton Power & Light Co.

Managing Director, Engineering Innovations Pte. Ltd. (Singapore)

President, Hathaway Automation Technologies, Inc.

5. My current position is Consultant, Information and Communication Technologies.
6. I am not an employee of Appellant or its assignees nor do I have a financial interest dependent upon the outcome of the above referenced application. However, I am being compensated for submitting this declaration, and I may have a financial interest dependent on application(s) that are either directly related or indirectly related to the above referenced application.

INDEPENDENT CLAIM 1

7. I have reviewed Applicant's Independent Claim 1, which as currently presented recites as follows:
 1. A method comprising:
 - (a) aggregating at least a part of one or more mote-addressed content indexes from a first set of motes administered by a first network administrator owned or controlled by a first business entity to form a first aggregated mote-addressed content index; and
 - (b) exposing at least a part of the first aggregated mote-addressed content index to an aggregator of (i) a first-set content index from the first set of motes administered by the first network administrator owned or controlled by the first business entity and (ii) a second-set content index from a second set of motes administered by a second network administrator owned or controlled by a second business entity.
8. I have reviewed the following portions of Mulgund et al. (U.S Application Number 2002/0161751) ("Mulgund"), "TAG: a Tiny Aggregation Service for Ad-Hoc Sensor Networks" by Samuel Madden et. al. ("Madden TAG"), and Simon et al. (US Patent No. 7,665,126 B2) ("Simon") cited by the USPTO with respect to Applicant's Independent Claim 1:

Method of and system for aggregating into a relational database model the state of an ad hoc network comprised of uniquely addressable distributed sensor nodes in communication using networking protocols with one another through links and to a database server through access points. A relational database logical design resident on the database server is dynamically updated with respect to the sensor network's current and historical topological information through the use of a traversal and interrogating network modeling agent. The distributed sensors nodes may be mobile, and may communicate by wired or wireless means through networking protocols such as the Internet.

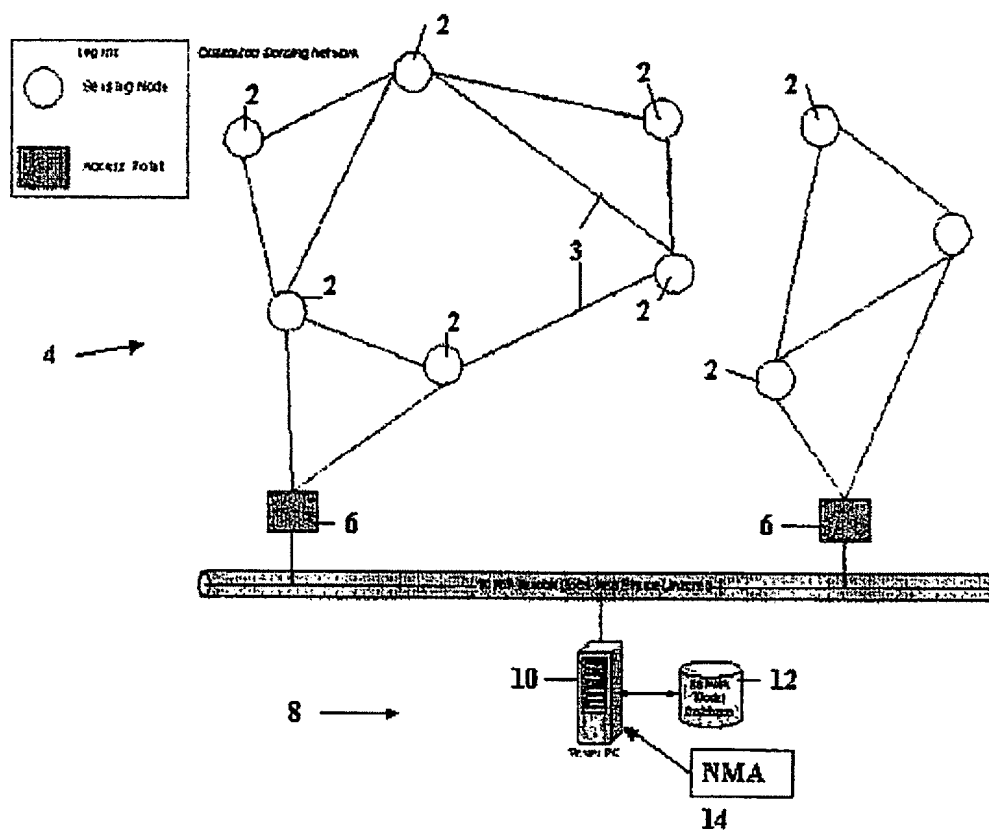
See *Mulgund* (Abstract)

The tools needed to implement the vision of seamless, global access to remote information are available only in part, and not yet as an integrated package. The Applicants describe below the development of an information architecture, which is referred to in certain embodiments as Intelemetric™, and a method of using the architecture which make it possible to aggregate, store, process, and distributed, real-time distributed sensor data into the enterprise, and make resulting information readily available over the Internet.


See *Mulgund* (paragraph [0005])

It is of no concern how this network topology came into being, how it is organized, what routing algorithms are used to pass messages from one node to the next, but rather, how to aggregate the information at each of the nodes into an off-network repository or network model database 12. The sensing nodes 2 may be mobile, and the interconnections may change over time. Furthermore, new nodes may join the network 4 at any time, and existing nodes may leave the network unexpectedly.

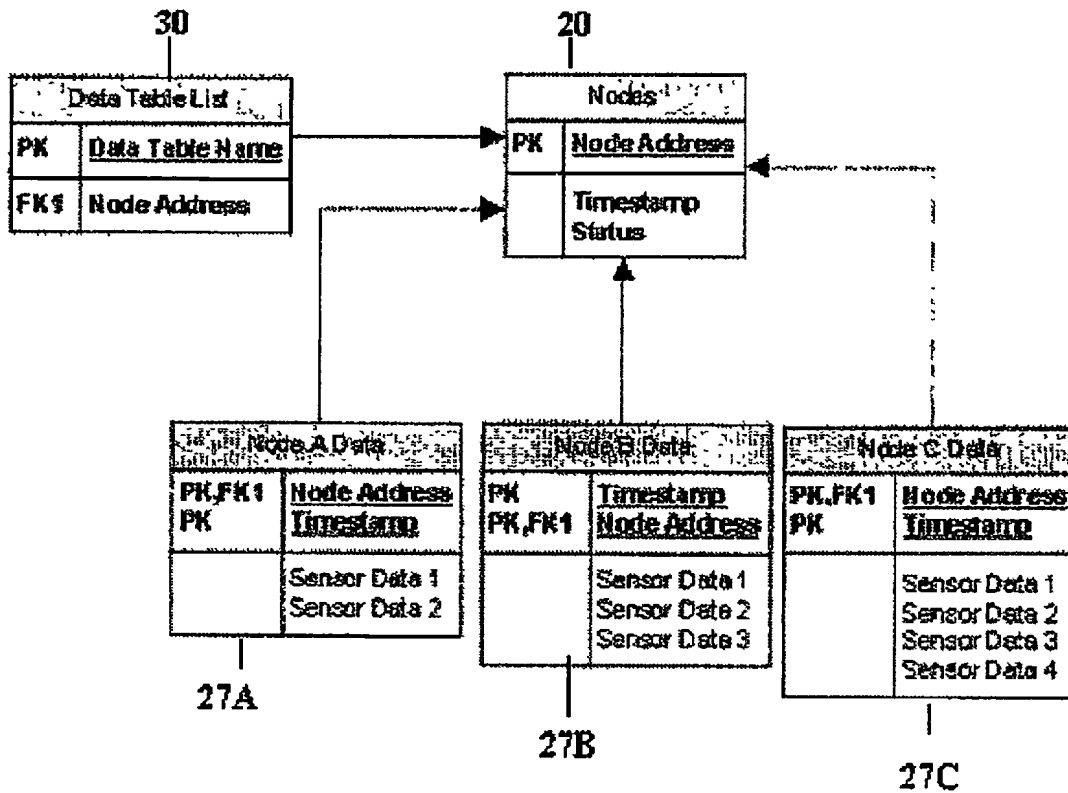
See *Mulgund* (paragraph [0025])



See Mulgund (Fig. 1)



See Mulgund (Fig. 3)



See Mulgund (Fig. 4)

We present the Tiny AGgregation (TAG) service for aggregation in TinyOS. TAG allows users to express simple, declarative queries and have them distributed and executed efficiently in networks of low-power, wireless sensors. We discuss various generic properties of aggregates, and show how those properties affect the performance of our in-network approach. We include a performance study demonstrating the advantages of our approach over traditional centralized, out-of-network methods, and discuss a variety of optimizations for improving the performance and fault-tolerance of the basic solution.

TAG operates as follows: users pose aggregation queries from a powered, storage-rich base station. Operators that implement the query are distributed into the network by piggybacking on the existing ad hoc networking protocol. Sensors route data back towards the user through a routing tree rooted at the base station. As data flows up this tree, it is aggregated according to an

aggregation function and value-based partitioning specified in the query. For example, consider the problem of counting the number of nodes in a network of indeterminate size. First, the request to count is injected into the network. Then, each leaf node in the tree reports a count of 1 to their parent; interior nodes sum the count of their children, add 1 to it, and report that value to their parent. Counts propagate up the tree in this manner, and flow out at the root.

Given the simple routing protocol from Section 2.1 and our SQL-like query model, we now discuss the implementation of the core TAG algorithm for in-network aggregation.

A naive implementation of sensor network aggregation would be to use a centralized, server-based approach where all sensor readings are sent to the base station, which then computes the aggregates. In TAG, however, we compute aggregates in-network whenever possible, because, if properly implemented, this approach can be lower in number of message transmissions, latency, and power consumption than the server-based approach. We will measure the advantage of in-network aggregation in Section 5 below; first, we present the basic algorithm in detail. We first consider the operation of the basic approach in the absence of grouping; we show how to extend it with grouping in Section 4.2.

4.1 Tiny Aggregation

TAG consists of two phases: a distribution phase, in which aggregate queries are pushed down into the network, and a collection phase, where the aggregate values are continually routed up from children to parents. Recall that our query semantics partition time into epochs of duration, and that we must produce a single aggregate value (when not grouping) that combines the readings of all sensors in the network during that epoch.

Given our goal of using as few messages as possible, the collection phase must insure that parents in the routing tree wait until they have heard from their children before propagating an aggregate value for the current epoch. We will accomplish this by having parents subdivide the epoch such that children are required to deliver their partial state records during a parent-specified time interval. This interval is selected such that there is enough time for the parent to combine partial state records and propagate its own record to its parent.

4.2 Grouping

Grouping in TAG is functionally equivalent to the GROUP BY clause in SQL: each sensor reading is placed into exactly one group, and groups are partitioned according to an expression over one or more attributes. The basic grouping technique is to push the expression down with the query, ask sensors to choose the group they belong to, and then, as answers flow back, update aggregate values in the appropriate groups.

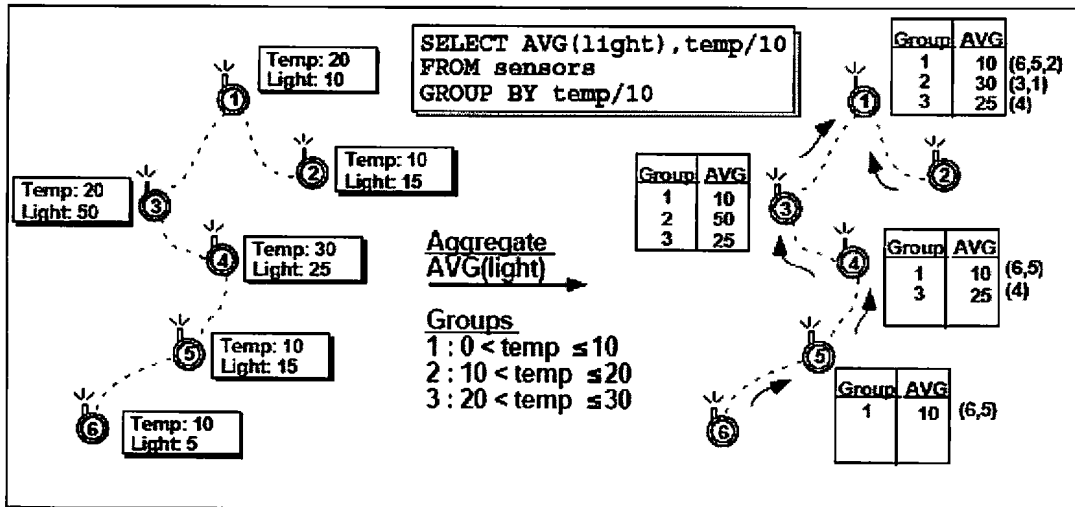
Partial state records are aggregated just as in the approach described above, except that those records are now tagged with a group id. When a sensor is a leaf, it applies the grouping expression to compute a group id. It then tags its partial state record with the group and forwards it on to its parent. When a sensor receives an aggregate from a child, it checks the group id. If the child is in the same group as the sensor, it combines the two values using the combining function. If it is in a different group, it stores the value of the child's group along with its own value for forwarding in the next epoch. If another child message arrives with a value in either group, the sensor updates the appropriate aggregate. During the next epoch, the sensor will send out the value of all the groups it collected information about during the previous interval, combining information about multiple groups into a single message as long as the message size permits. Figure 2 shows an example of computing a query grouped by temperature that selects average light readings.

Recall that queries can also contain a HAVING clause, which constrains the set of groups in the final query result. We sometimes pass this predicate into the network along with the grouping expression. The predicate is only sent into the network if it can potentially be used to reduce the number of messages that must be sent: for example, if the predicate is of the form $\text{MAX}(\text{attr}) < x$, then information about groups with $\text{MAX}(\text{attr}) > x$ need not be transmitted up the tree, and so the predicate is sent down into the network. When a node detects that a group does not satisfy a HAVING clause, it can notify other nodes in the network of this information to suppress transmission and storage of values from that group. Note that HAVING clauses can be pushed down only for monotonic aggregates; non-monotonic aggregates are not amenable to this technique. However, not all HAVING predicates on monotonic aggregates can be pushed down; for example, $\text{MAX}(\text{attr}) > x$, cannot be applied in the network because a node cannot know that, just because its local value of x is less than the MAX over the entire group is less than x .

Because the number of groups can exceed available storage on any one (non-leaf) sensor, a way to evict groups is needed. Once an eviction victim is selected, it is forwarded to the sensor's parent, which may choose to hold on to

the group or continue to forward it up the tree. Notice that a single sensor may evict several groups in a single epoch (or the same group multiple times, if a bad victim is selected). This is because, once group storage is full, if only one group is evicted at a time, a new eviction decision must be made every time a value representing an unknown or previously evicted group arrives. Because groups can be evicted, the base station at the top of the network may be called upon to combine partial groups to form an accurate aggregate value. Evicting partially computed groups is known as partial preaggregation, as described in [11].

Thus, we have shown how to partition sensor readings into a number of groups and properly compute aggregates over those groups, even when the amount of group information exceeds available storage in any one sensor. We will discuss experiments with grouping and group eviction policies in Section 5.2. First, we summarize some of the additional benefits of TAG.



See *Madden TAG* (abstract, section 1.1 par. 2, section 4, 4.1 pars. 1-2, and 4.2; Fig. 2)

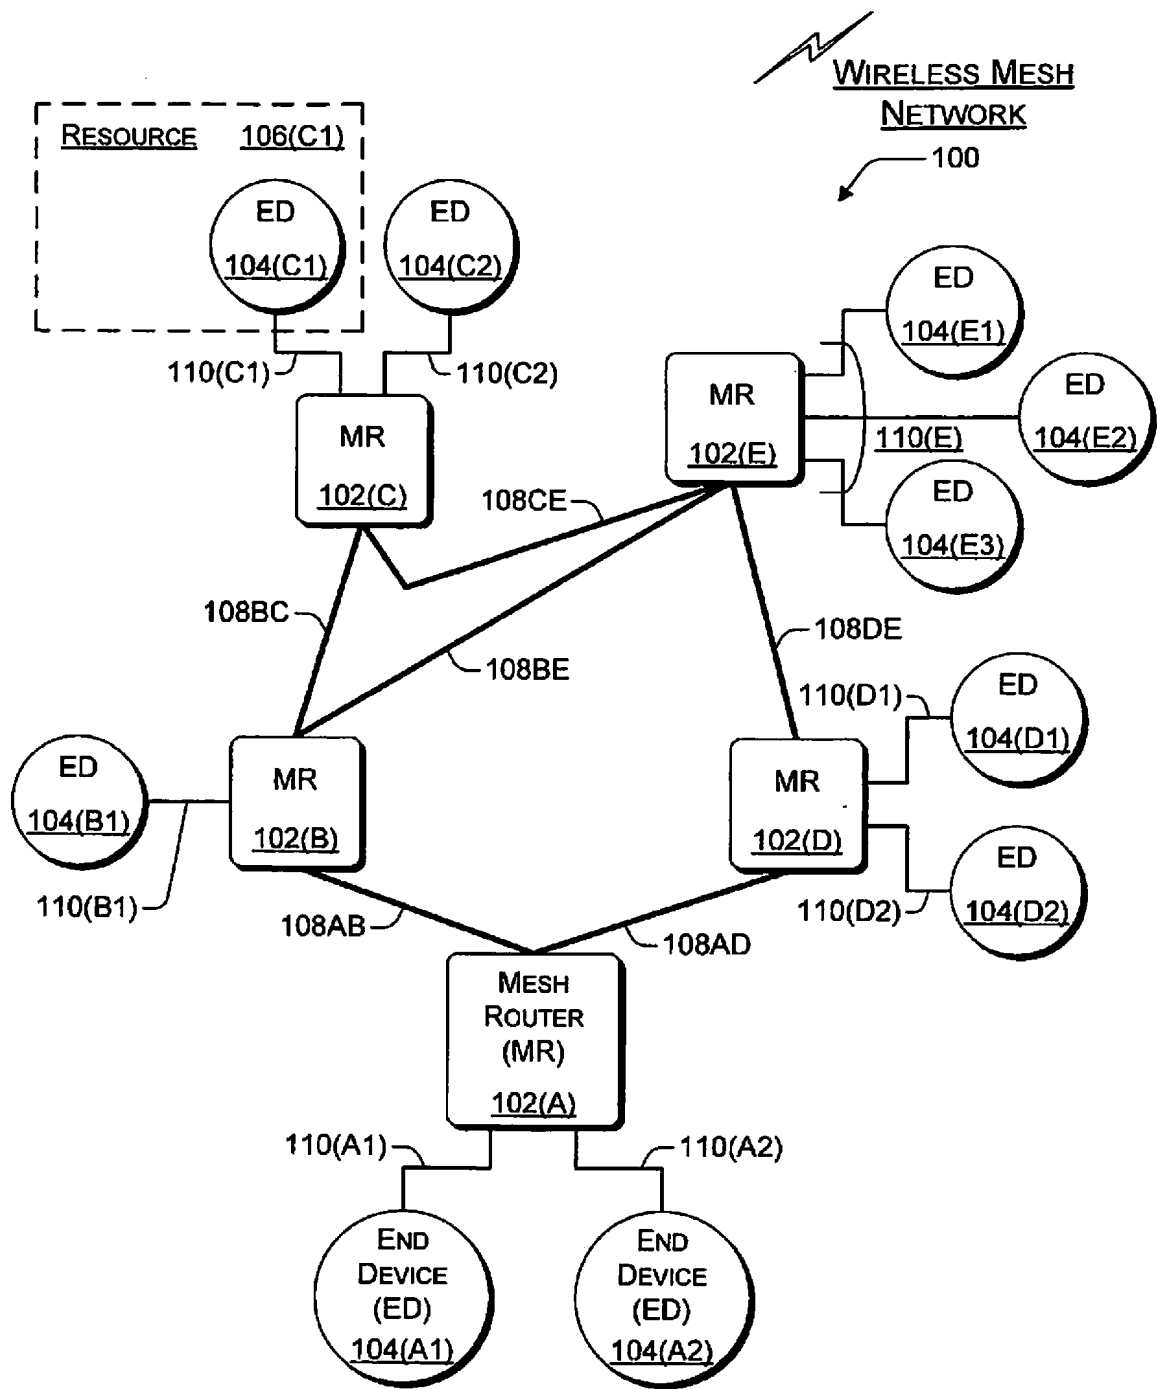


Fig. 1

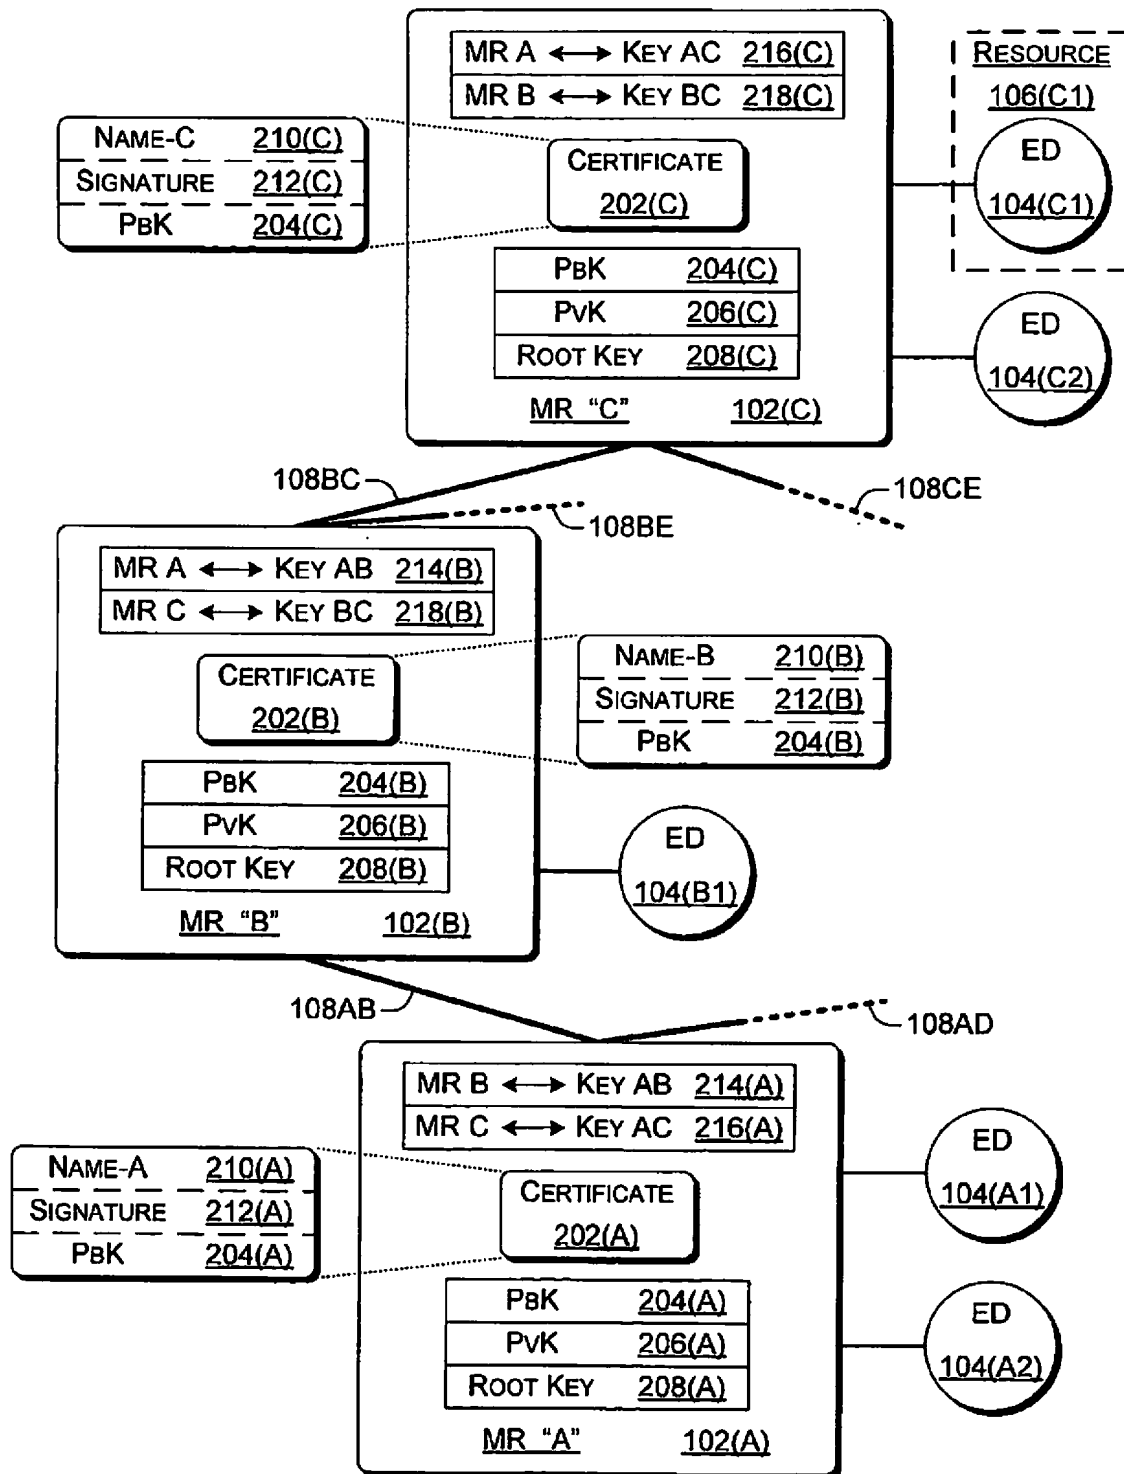


Fig. 2

In an exemplary method implementation, a method includes: designating a neighborhood administrator; receiving notification of a delinquent router from the designated neighborhood administrator; and excluding the delinquent router responsive to the notification. In an exemplary mesh router implementation, a mesh router is capable of establishing a wireless mesh network with other mesh routers, the mesh router is further capable of designating a neighborhood administrator mesh router; and the mesh router is adapted to exclude another mesh router that is associated with a particular certificate when the particular certificate has been identified as delinquent by the designated neighborhood administrator mesh router.

Although each mesh router 102 is illustrated as being in wireless communication with from one to three end devices 104, each may alternatively be in communication with any number of end devices 104. Mesh router 102(A) is in wireless communication with two end devices 104(A1) and 104(A2) over wireless links 110(A1) and 110(A2), respectively. Mesh router 102(B) is in wireless communication with one end device 104(B1) over wireless link 110(B1). Mesh router 102 (C) is in wireless communication with two end devices 104 (C1) and 104(C2) over wireless links 110(C1) and 110(C2), respectively. Similarly, mesh router 102(E) has wireless links 110(E) with three end devices 104(E1), 104(E2), and 104 (E3). Mesh router 102(D) has wireless links 110(D1) and 110(D2) to two end devices 104(D1) and 104(D2), respectively.

FIG. 2 illustrates an exemplary public key infrastructure (PKI) at the mesh router tier in which each mesh router 102 is associated with a certificate 202. Three exemplary mesh routers 102(A), 102(B), and 102(C) are specifically shown. As illustrated, each mesh router 102 includes a certificate 202, a public key (PbK) 204, a private key (PvK) 206, and a root key 208. Each certificate 202 includes a name 210, a signature 212, and the corresponding public key 204. Mesh router "A" 102(A) is used in particular to describe these general aspects of the exemplary PKI at the mesh router tier.

See *Simon* (Figs. 1 and 2; Abstract, col. 3 lines 4-18; col. 4 lines 23-32)

9. The USPTO may be alleging¹ that Independent Claim 1 can be mapped to the USPTO cited portions of the Reference(s) allegedly related to Independent Claim 1. If such is the

¹ As shown, the office action of the USPTO merely cites locations in the reference in parentheses which are not explained. The USPTO has NOT provided any claim interpretation nor a reasoned application of the cited reference(s) to the claim. As there is no mapping of the claim language to any specific aspects of the excerpt, it is unclear whether the examiner is alleging an interpretation of the claim that reaches the cited art, or is alleging some "teaching" of the cited art that reaches the claim language, or is alleging some combination of the foregoing. In the

case, in my expert opinion, the Reference(s) fail to recite several express recitations of Applicant's Independent Claim 1 in that, as shown following, the mapping would fail to take such express recitations into account.

a. Applicant's Independent Claim 1 as currently presented recites:

1. A method comprising:

(a) aggregating at least a part of one or more mote-addressed content indexes from a first set of motes administered by a first network administrator owned or controlled by a first business entity to form a first aggregated mote-addressed content index; and

(b) exposing at least a part of the first aggregated mote-addressed content index to an aggregator of (i) a first-set content index from the first set of motes administered by the first network administrator owned or controlled by the first business entity and (ii) a second-set content index from a second set of motes administered by a second network administrator owned or controlled by a second business entity.

b. With respect to Claim 1, the Action states as follows:

“As to claim 1, Mulgund teaches:

aggregating at least a part of one or more mote-addressed content indexes from a first set of motes [aggregating indexing information related to sensor data outputs stored in Sensor Data Table 24. It is well known to one of ordinary skill in the art that such indexing data allows one to distinguish between different sensor data outputs, as discussed in connection with relational databases of Figs. 3 and 4] (abstract, par. [0005] and [0025], Fig. 3, Fig. 4), wherein the first set of motes [nodes to the left of Fig. 1] is administered by a first network administrator {first network access point to the left of Fig. 1] and wherein the second set of motes [nodes to the right of Fig. 1] is administered by a second network administrator [second network access point to the right of Fig. 1].

Mulgund does not teach that said aggregating is performed in at least one mote in a second set of motes. Mulgund is silent as to whether the first and second network access points are owned or controlled by different legal entities.

Madden teaches aggregating at least a part of one or more mote addressed indexes [sensor attributes, such as group id] from the first set of motes as performed in at least one mote [parent node] in a second set of motes [a collection

absence of knowing, and in view of the statutory deadline for responding to rejections and the lack of a practicable mechanism to force the examiner to render a proper rejection, this declaration relates to the examiner's rejections in the alternative.

phase, where aggregate values are continually routed up from children to parents] (abstract, section 1.1 par. 2, section 4, 4.1 pars. 1-2, and 4.2; Fig. 2).

It would have been obvious to one of ordinary skill in the art at the time of the invention to modify the method of Mulgund by aggregating said indexes at a mote in order to lower the number of message transmissions, latency, and power consumption than the server-based approach (as taught by Mulgund) (Madden, section 4 under In- Network Aggregates).

Simon is directed to a method of controlling access of mesh routers to a network resources based on information contained in a certificate associated with the particular router. See abstract. Simon teaches the first set of devices (104C) is administered by a first network administrator (102C) controlled by a first legal entity [mesh router 102C is controlled by the associated certificate (202C)] and wherein the second set of devices (104B) is administered by a second network administrator (102B) controlled by a second legal entity (mesh router 102B is controlled by the associated certificate (202B)) (Figs. 1 and 2; col. 3 lines 4-18; col. 4 lines 23-32).

It would have been obvious to one of ordinary skill in the art at the time of the invention to modify the method of Mulgund in view of Madden by having the first and second network access points being owned or controlled by different legal entities in order to introduce control and/or accountability into spontaneously-formed wireless networks (col. 1 lines 43-45 of Simon)."

Examiner's *Office Action*, pp. 18-20 (4 March 2010)

The USPTO may be alleging² that Independent Claim 1 can be mapped to the USPTO cited portions of the Reference(s), allegedly related to Independent Claim 1, as follows: It appears, in view of USPTO prior rejection, to Declarant that the USPTO may try to map "(b) **exposing at least a part of the first aggregated mote-addressed content index to an aggregator of (i) a first-set content index from the first set of motes administered by the first network administrator owned or controlled by the first business entity and (ii) a**

² As shown, the office action of the USPTO merely cites locations in the reference in parentheses which are not explained. The USPTO has not provided claim interpretation nor a reasoned application of the cited reference(s) to the claim. Parenthetical indications of art placed in proximity to claim language is unclear since it is impossible to know whether the examiner is alleging an interpretation of the claim that reaches the cited art, or is alleging some "teaching" of the cited art that reaches the claim language, or is alleging some combination of the foregoing. In the absence of knowing, and in view of the statutory deadline for responding to rejections and the lack of a practicable mechanism to force the examiner to render a proper rejection, declarant herein analyzes the examiner's rejections in the alternative.

second-set content index from a second set of **motes administered by a second network administrator owned or controlled by a second business entity**” (Emphasis added) onto a configuration in Simon by which “an exemplary public key infrastructure (PKI) at the mesh router tier in which each mesh router 102 is associated with a certificate 202.” It appears that the USPTO has not yet explained how it could reach such mappings under the framework of the broadest reasonable interpretation consistent with the specification as is the USPTO’s burden (e.g., such as by examples drawn from Applicant’s claims or detailed description), and furthermore, Declarant points out that this mapping would not address at least the “exposing ... content index to an aggregator of (i) ... motes administered by the first network administrator owned or controlled by the first business entity and (ii) ... motes administered by a second network administrator owned or controlled by a second business entity.”

- c. Based upon my education and/or experience, my opinion is that any mapping of the above recitations of Independent Claim 1 onto the cited portions of the Reference(s) would be unreasonable because such a mapping would overlook express recitations of Independent Claim 1. For example, claim 1 recites “a first-set content index from the first set of motes administered by the first network administrator owned or controlled by the first business entity” and “a second-set content index from a second set of motes administered by a second network administrator owned or controlled by a second business entity”, and these recitations can be considered in view of the following listed examples set forth in the specification (e.g., detailed description):

- i. The Detailed Description recites:

- 1. “In some implementations, first administered set 800 of motes constitutes all or part of a network under a first administrator and second-administered set 802 of motes constitutes all or part of a network under a second administrator, where the first and/or second administrators tend not to have any significant knowledge

of the internal operations of networks they don't administer. Examples in which this may be the case are where first administered set 800 and second-administered set 802 are owned by different business entities, and where first-administered set 800 and second-administered set 802 have been constructed for two separate purposes (e.g., one set to monitor crops and the other set to monitor building systems, and thus the systems were not designed to interact with each other).”

See U.S. Patent Application Publication No. 2005/0220142, paragraph [0075]”

d. As noted above, the USPTO has yet to take into account at least the “**exposing at least a part of the first aggregated mote-addressed content index to an aggregator of (i) a first-set content index from the first set of notes administered by the first network administrator owned or controlled by the first business entity and (ii) a second-set content index from a second set of notes administered by a second network administrator owned or controlled by a second business entity**” recitations of claim 1. Consequently, the Action cites no specific art in relation to such claim recitations, as is required under 37 C.F.R. 1.104(c) (2). In view of the lack of USPTO-cited evidence for at least these yet to be addressed express claim recitations, I conclude, based on my education and/or experience, that the cited art does not teach Independent Claim 1.

10. The USPTO may be alleging³ that the USPTO-cited portions of the Reference(s) “teach” Independent Claim 1. If such is the case, in my expert opinion, on the record evidence

³ As shown, the office action of the USPTO merely cites locations in the reference in parentheses which are not explained. The USPTO has not provided claim interpretation nor a reasoned application of the cited reference(s). Parenthetical indications of art placed in proximity to claim language is unclear since it is impossible to know whether the examiner is alleging an interpretation of the claim that reaches the cited art, or is alleging some “teaching” of the cited art that reaches the claim language, or is alleging some combination of the foregoing. In the absence of knowing, and in view of the statutory deadline for responding to rejections and the lack of a practicable mechanism to force the examiner to render a proper rejection, declarant herein analyzes the examiner’s rejections in the alternative.

specifically cited by the USPTO, the Reference(s) cannot be characterized to teach the text of Applicant's Independent Claim 1 for the reasons shown following.

a. Applicant's Independent Claim 1 recites:

“1. A method comprising:

(a) aggregating at least a part of one or more mote-addressed content indexes from a first set of motes administered by a first network administrator owned or controlled by a first business entity to form a first aggregated mote-addressed content index; and

(b) **exposing** at least a part of the first aggregated mote-addressed content index to an aggregator of (i) a first-set content index from the first set of motes administered by the first network administrator owned or controlled by the first business entity and (ii) a second-set content index from a second set of motes administered by a second network administrator owned or controlled by a second business entity.” (Emphasis added)

b. With respect to Claim 1, the Action states as follows:

Madden teaches aggregating at least a part of one or more mote addressed indexes [sensor attributes, such as group id] from the first set of motes as performed in at least one mote [parent node] in a second set of motes [a collection phase, where aggregate values are continually routed up from children to parents] (abstract, section 1.1 par. 2, section 4, 4.1 pars. 1-2, and 4.2; Fig. 2).

It would have been obvious to one of ordinary skill in the art at the time of the invention to modify the method of Mulgund by aggregating said indexes at a mote in order to lower the number of message transmissions, latency, and power consumption than the server-based approach (as taught by Mulgund) (Madden, section 4 under In- Network Aggregates).

Simon is directed to a method of controlling access of mesh routers to a network resources based on information contained in a certificate associated with the particular router. See abstract. Simon teaches the first

set of devices (104C) is administered by a first network administrator (102C) controlled by a first legal entity [mesh router 102C is controlled by the associated certificate (202C)] and wherein the second set of devices (104B) is administered by a second network administrator (102B) controlled by a second legal entity (mesh router 102B is controlled by the associated certificate (202B)) (Figs. 1 and 2; col. 3 lines 4-18; col. 4 lines 23-32).

- c. The USPTO cited portion of the Reference(s), which the PTO may allege as being related to herein amended Independent Claim 1, or allegedly related to unamended Independent Claim 1, recites (or directly discloses):

“TAG consists of two phases: a distribution phase, in which aggregate queries are pushed down into the network, and a collection phase, where the aggregate values are continually routed up from children to parents. Recall that our query semantics partition time into epochs of duration, and that we must produce a single aggregate value (when not grouping) that combines the readings of all sensors in the network during that epoch.”

See *Madden TAG* (section 4.1)

“When a sensor receives an aggregate from a child, it checks the group id. If the child is in the same group as the sensor, it combines the two values using the combining function. If it is in a different group, it stores the value of the child’s group along with its own value for forwarding in the next epoch. If another child message arrives with a value in either group, the sensor updates the appropriate aggregate.”

See *Madden TAG* (section 4.2)

“In an exemplary method implementation, a method includes: designating a neighborhood administrator; receiving notification of a delinquent router from the designated neighborhood administrator; and excluding the delinquent router responsive to the notification. In an exemplary mesh router implementation, a mesh router is capable of establishing a wireless mesh network with other mesh routers, the mesh router is further capable of designating a neighborhood administrator mesh router; and the mesh router is adapted to exclude another mesh router that is

associated with a particular certificate when the particular certificate has been identified as delinquent by the designated neighborhood administrator mesh router.”

See *Simon* (Abstract)

“An exemplary public key infrastructure (PKI) at the mesh router tier in which each mesh router 102 is associated with a certificate 202.”

See *Simon* (col. 4 lines 23-32)

d. Based upon my education and experience, my opinion is that the cited portions of the Reference(s) are, at first sight, on the first appearance, on the face of it, so far as can be judged from the first disclosure, very different on their faces in that the USPTO-identified portions of *Simon* and *Madden TAG* do not recite the text of at least amended Clause [b] of Independent Claim 1: “exposing at least a part of the first aggregated mote-addressed content index to an aggregator of (i) a first-set content index from the first set of motes administered by the first network administrator owned or controlled by the first business entity and (ii) a second-set content index from a second set of motes administered by a second network administrator owned or controlled by a second business entity.” Instead, *Simon* directly recites that (1) “the mesh router is further capable of designating a neighborhood administrator mesh router; and the mesh router is adapted to exclude another mesh router that is associated with a particular certificate when the particular certificate has been identified as delinquent by the designated neighborhood administrator mesh router” and (2) “an exemplary public key infrastructure (PKI) at the mesh router tier in which each mesh router 102 is associated with a certificate 202.”

Based upon my education and experience, my opinion is that the above recitation of clause [b] of Independent Claim 1 and the cited portions of the Reference(s) allegedly related to unamended clause [b] are, at first sight, on the first appearance, on the face of it, so far as can be judged from the first disclosure, very different on their faces. The USPTO characterizes the cited portions of the References to teach unamended clause [b] but such characterization are not supported. Also as

noted above, the USPTO has asserted that the Simon reference “is directed to a method of controlling access of mesh routers to a network resources based on information contained in a certificate associated with the particular router” and that “Simon teaches the first set of devices (104C) is administered by a first network administrator (102C) controlled by a first legal entity [mesh router 102C is controlled by the associated certificate (202C)] and wherein the second set of devices (104B) is administered by a second network administrator (102B) controlled by a second legal entity (mesh router 102B is controlled by the associated certificate (202B))”. Declarant disagrees with this assertion and notes that the Simon reference actually recites that “the mesh router is adapted to exclude another mesh router that is associated with a particular certificate” and “an exemplary public key infrastructure (PKI) at the mesh router tier in which each mesh router 102 is associated with a certificate 202.” To Declarant, it appears that the USPTO may try to close a significant gap between these actual recitations of the Simon reference and the structure of “[b] exposing at least a part of the first aggregated mote-addressed content index to an aggregator of (i) a first-set content index from the first set of motes administered by the first network administrator owned or controlled by the first business entity and (ii) a second-set content index from a second set of motes administered by a second network administrator owned or controlled by a second business entity.” (Emphasis Added) (in Clause [b] of Applicant’s Claim 1) without providing any evidence, by merely making an unsupported assertion.

To Declarant, it appears that the USPTO may try to equate the teaching of the Simon reference and the structure of amended claim 1. The Simon reference recites that a “mesh router is adapted to exclude another mesh router that is associated with a particular certificate” and “an exemplary public key infrastructure (PKI) at the mesh router tier in which each mesh router 102 is associated with a certificate 202.” The structure of clause [b] of amended claim 1 recites “[b] **exposing** at least a part of **the first aggregated mote-addressed content index to an aggregator of (i) a first-set content index from the first set of motes administered by the first network**

administrator owned or controlled by the first business entity and (ii) a second-set content index from a second set of **motes administered by a second network administrator owned or controlled by a second business entity.**” (Emphasis Added) To Declarant, it appears that the USPTO may try to close a significant gap between these actual recitations of the Simon reference and the structure of “[b] exposing at least a part of the first aggregated mote-addressed content index to an aggregator of (i) a first-set content index from the first set of motes administered by the first network administrator owned or controlled by the first business entity and (ii) a second-set content index from a second set of motes administered by a second network administrator owned or controlled by a second business entity.” (Emphasis Added) (in Clause [b] of Applicant’s Claim 1) without providing any evidence, by merely making an unsupported assertion.

Based upon my education and experience, my opinion is that the cited portions of the Reference(s) are, at first sight, on the first appearance, on the face of it, so far as can be judged from the first disclosure, very different on their faces in that the USPTO-identified portions Madden TAG do not recite the text of at least amended Clause [b] of Independent Claim 1: “**exposing** at least a part of **the first aggregated mote-addressed content index to an aggregator of** (i) a first-set content index from the first set of **motes administered by the first network administrator owned or controlled by the first business entity** and (ii) a second-set content index from a second set of **motes administered by a second network administrator owned or controlled by a second business entity.**” (Emphasis added) The cited portions of Madden TAG recite that “TAG operates as follows: users pose aggregation queries from a powered, storage-rich base station. Operators that implement the query are distributed into the network by piggybacking on the existing ad hoc networking protocol. Sensors route data back towards the user through a routing tree rooted at the base station”. To Declarant, it appears that the USPTO may try to close a significant gap between these actual recitations of the Madden TAG reference and the structure of “[b] exposing at least a part of the first aggregated

mote-addressed content index to an aggregator of (i) a first-set content index from the first set of motes administered by the first network administrator owned or controlled by the first business entity and (ii) a second-set content index from a second set of motes administered by a second network administrator owned or controlled by a second business entity.” (Emphasis Added) (in Clause [b] of Applicant’s Claim 1) without providing any evidence, by merely making an unsupported assertion.

In summary to Declarant, it appears that the USPTO may try to close a significant gap between the amended clause [b] and the actual recitations of the Madden TAG and Simon reference as shown above⁴ as the USPTO-identified portions of Simon and Madden TAG References do not recite the text of at least amended Clause [b] of Independent Claim 1: “exposing at least a part of the first aggregated mote-addressed content index to an aggregator of (i) a first-set content index from the first set of motes administered by the first network administrator owned or controlled by the first business entity and (ii) a second-set content index from a second set of motes administered by a second network administrator owned or controlled by a second business entity.”⁵

- e. The Action provides neither convincing evidence nor convincing syllogistic arguments to resolve at least the noted facially apparent differences between the direct disclosures of the cited portions of the Reference(s) allegedly related to clause [a] of herein amended Independent Claim 1 and clause [b] of herein amended Independent Claim 1; indeed, Examiner does not even note such facially apparent differences in his Office Action. Accordingly, on the record evidence cited by Examiner, and as one skilled in the art, I conclude that the cited portions of the Reference(s) allegedly related to clause [a] and/or [b] of herein amended Independent Claim 1 do not teach herein amended Independent Claim 1 for at

⁴ Neither do the USPTO-identified portions of Simon recite “aggregating at least a part of one or more mote-addressed content indexes from a first set of motes administered by a first network administrator owned or controlled by a first business entity to form a first aggregated mote-addressed content index,” as recited in Clause [a].

⁵ The USPTO has stated “Mulgund is silent as to whether the first and second network access points are owned or controlled by different legal entities.”

least the foregoing reasons. Specifically, in view of such lack of objective evidence or convincing syllogistic arguments resolving such facial differences, my expert opinion is that the examiner-cited art does not teach Independent Claim 1 for at least the foregoing reasons.

- f. Moreover, not only does the Reference(s) not teach Independent Claim 1, but in view of the foregoing noted facially apparent differences between the direct disclosures of the cited portions of the Reference(s) allegedly related to clause [a] and/or [b] and the recitations of herein amended Independent Claim 1, and in view of the record evidence cited by Examiner, my opinion is that there is no reasonable basis for modifying the identified one or more of the technologies of Mulgund⁶ in view of Madden TAG and/or Simon allegedly related to clause [a] and/or [b] to meet the recitations of herein amended Independent Claim 1.

11. I hereby further declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Richard Lord

Richard Lord

8/9/2010

Date

⁶ Neither do the USPTO-identified portions of Mulgund disclose "a first set of notes ..." as recited in Clause [a] of claim 1. In view of applicant's specification, I am unable to equate the "uniquely addressable distributed sensor nodes" as recited in the Mulgund reference with "a first set of notes" as recited in claim 1.